Proposal for Validation of a Multispectral Synthetic Scene Generator Model

by:

Marcos C. Sola, Mark W. Orletsky, Quochien B. Vuong, and Charles R. Kohler

Synthetic Imagery Branch, Sensors and Electron Devices Directorate,
U. S. Army Research Laboratory, 2800 Powder Mill Road, Adelphi, Maryland 20783.

ABSTRACT

The establishment of a sufficient, field-measured database to support the analysis of ATR algorithms, sensor fusion effectiveness, and sensor system performance for multiple combinations of targets, environments, sensors, and locations will severely challenge the limited, available resources within the Army. However, the use of a high-resolution, synthetic scene generator model (SSGM) for time-independent applications can alleviate the database requirement. A methodology for a robust validation of SSGM is proposed, which will consist of defining sets of images (real and corresponding SSGM imageries) and using human observers to define a baseline. First-order comparisons of a real scene to a synthetic scene will be performed with the use of the filters in the TARDEC [20] model or a comparable computational vision model (CVM). The similarity of target to background histograms as a function of various CVM filters will need to be analyzed to define first order-effects. Second-order metrics are defined in terms of probability of detection, detection timeline, and false alarm rate. A metric for the target signature will be mathematically defined to test these second-order effects. For a given application. the necessary and sufficient metrics are discussed.

1. INTRODUCTION

A great number of Army programs require a significant number of signature databases in order to satisfy program development. A partial listing of these programs include Intelligence Signature Assessment, Battlefield Visualization Test Bed (BVTB), ATR Algorithm Development, Sensor Performance Analyses, Multi-Sensor Fusion, Countermeasure/Counter-Countermeasure (CM/CCM), Target Acquisition (TA) Modeling Improvement, Required Operational Capability Requirements, Tri-Service Smart Missile/Munitions Testing, and Computer War Gaming Input.

Any of the programs listed above would require signatures related to target acquisition. Target acquisition is a function of many variables; among them are the target, background, environment, geographical location, time, and sensor. Each of these sets is composed of subsets. To physically measure signatures in the field to satisfy a wide dynamic range of signature conditions would be a substantial budgetary challenge for any project manager [14]. The physically measured data must not only answer the requirements of the project, but also must be general and of sufficient resolution to take into account other near term signature requirements [21]. Hence, there is the need to make use of synthetic databases to augment, supplement, and/or complement field databases. However, this requires a robust validation of the particular synthetic scene generator model (SSGM) that is used so that it gains credibility and acceptance in the modeling and simulation community.

2 APPROACH

An integral step in the validation of an SSGM is to compare a set of image pairs under the criteria that are relevant to the given application. An image pair is defined here to be the physically measured scene (target in background/clutter) and the corresponding synthetically generated scene based on model input requirements. For example, if we had a real-synthetic pair of images of a simple block world, with well-defined lighting, angles, etc., we would almost definitely do well by comparing the images on a pixel-to-pixel basis. While not every pixel of the synthetic image would have the same value as the corresponding pixel of the real image, a sufficient number would be close enough to indicate the quality of the synthesis. In this case, the simplicity of the "world" under consideration would almost definitely permit synthesis of images that are pixel-wise very similar to real images.

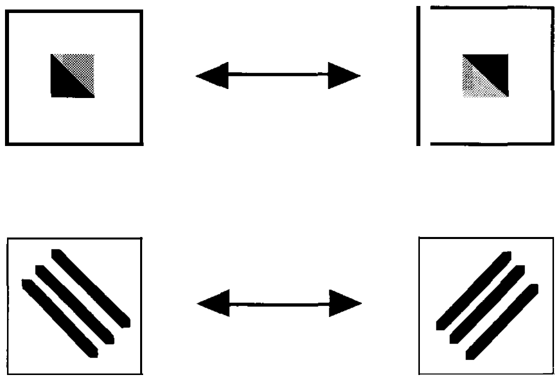

Figure 1. Examples of different images that have the same grey-level histograms.

For most applications of synthetic images, however, a bit-wise comparison to a corresponding real image is impractical as a validation criterion. Take, for example, the task of comparing images of a natural scene containing trees, grass, sky, water, etc. Here the variations that could be expected between the images would be large. For example, the wind may sway the objects in different directions, the clouds may cast shadows in different places, and so forth. Thus, the use of a pixel-to-pixel comparison would almost definitely be too exacting to effectively evaluate the quality of the synthesis.

An alternative method of comparison that has seen substantial use involves the comparison of statistics [2, 3, 4, 5, 8,9, 20] that have been derived over the entire image. Such statistics include grey-level histograms, local-energy histograms, and many others. These statistics frequently give information about the quality of a synthetic image, but rely on obtaining the statistics from the image as a whole, and thus can be misleading, as the following examples show.

Intensity histogram (also called grey-level histogram) is a necessary but not sufficient tool for comparative assessment of an image pair. Figure 1 shows two different image pairs that would produce the same intensity histogram for both images in each of the given pairs, while Figure 2 shows two different weapon platforms with exactly the same intensity histogram. Figure 2 was produced by a C-program that repositions the pixels of one image to approximate the other. Thus virtually any image can be slightly modified to have a grey-level histogram of another image while still retaining its original "look". Thus, given this example, sole reliance on histogram distribution as a similarity-metric for image comparison can lead to a wrong conclusion.

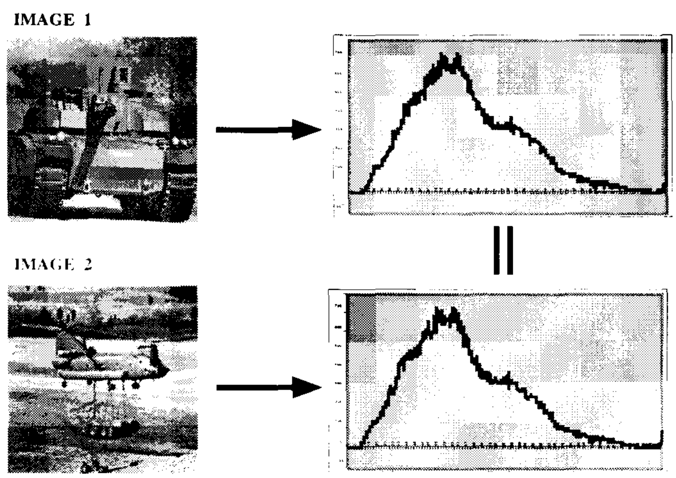

Figure 2. A tank and a helicopter with their identical grey-level histogram.

The validation approach taken by the Army Research Laboratory (ARL) is to develop a robust SSGM validation methodology for any synthetic rendering model. In particular, we would like to test this methodology against our own model for generating synthetic images, CREATION [6, 7, 10, 11, 17, 18]. Figure 3 is a generic approach to the problem of image validation. An "image pair" is passed through a comparator, and its output is statistically analyzed to identify a validation metric.

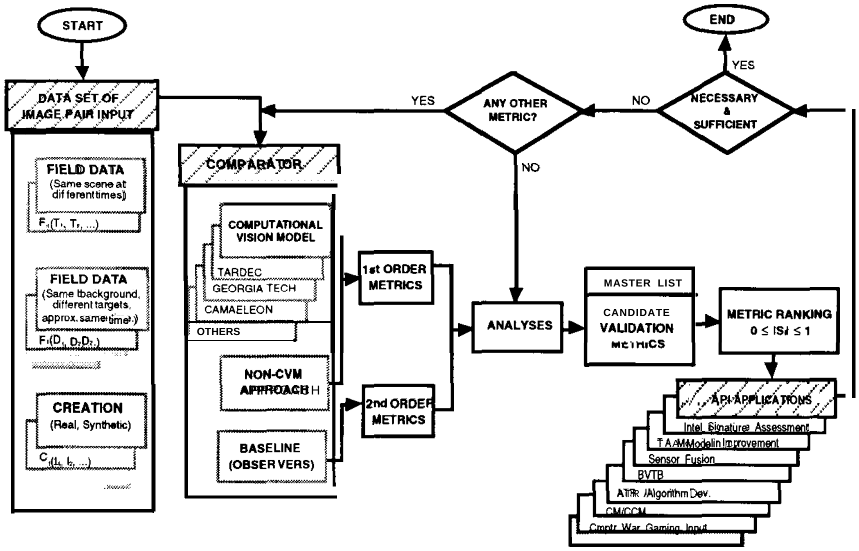

Figure 3. Diagram of an approach to image validation.

The comparator is composed of three components: a computation vision model, a non-computational vision model approach, and always a group of human observers used as a baseline. First-order validation metrics are obtained from a computational vision model such as the Georgia Tech Vision Model [2, 3], TARDEC [20] Vision Model, German CAMAELEON [4, 5, 9] Model, or other suitable candidates. Second-order metrics are obtained by using trained, human observers to provide probability of detection, highest level of acquisition (classification, recognition, identification) given detection, detection timeline, and false alarm rate. The noncomputational vision model approach can provide either first-order or second-order metrics. Note that if an image has the same second-order metrics as another, this does not necessarily mean that their validation metrics will be the same. Consider as an example two images that provide all the same second-order metrics as previously mentioned. Such a condition can be satisfied, for example, by a low observable tank at close range vs. a high contrast tank at long range (all other variables being equal). A robust set of validation metrics should provide an indication of two different scene conditions.

From the comparator, a master list of candidate validation metrics is compiled. A sufficient set of image pairs is parsed through the particular component of the comparator that is being tested for statistical analysis. These metrics are then "rank-ordered" from zero to one in terms of how similar the synthetic scene is compared to the real scene. A similarity of one is a perfect fit, while a similarity of zero signifies no correlation between the two images being compared. The number of validation metrics required is a function of the SSGM application. For high-resolution applications, such as target acquisition modeling improvement, most if not all the validation metrics with high similarity values may be required. For low-resolution applications, such as real time computer war gaming, only a certain portion of the validation metrics may be needed and their similarity metric requirement will be less stringent. It is therefore necessary to rank-order the validation metrics from the master list for a given synthetic scene-rendering application. In the CREATION model, we have a problem of field measured data from which we can generate a synthetic scene. Part of the problem is that this particular model requires a 24-hour diurnal cycle target signature history in order for us to be able to create a synthetic scene. Note that although ARL needs to validate the CREATION model, it must first be able to develop a robust validation methodology. To alleviate the image pair database problem, some pairs could be created from real field data, as, for example, scenes at different times or for the same time and background, but different targets.

3.0 PRESENT METHODOLOGY

Our current validation approach is to postulate a similarity metric that is defined as:

$$0 \le |S_i| \le 1, \qquad \text{where: } S_i = \begin{cases} 0, \text{ no match} \\ 1, \text{ perfect match} \end{cases}$$

and i = statistical validation metric (0, 1,2, ...n) in the master validation metric list.

Real **Image** **Synthetic Image**

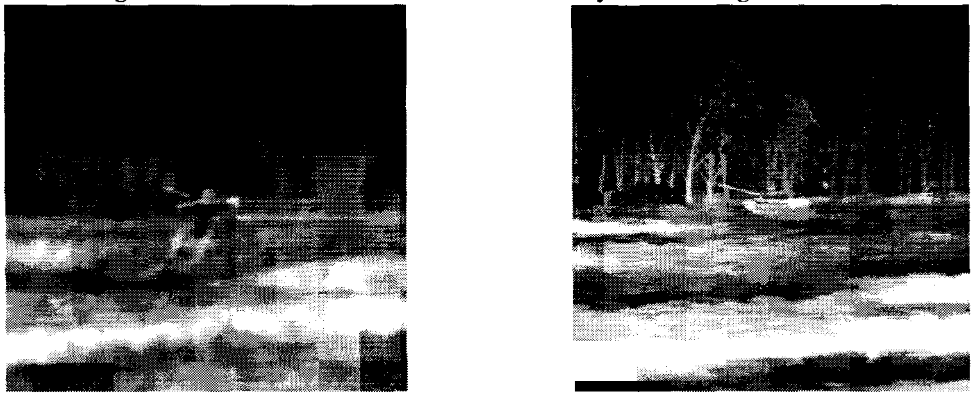

(a) (b)

Figure 4. M60-A1 tank.

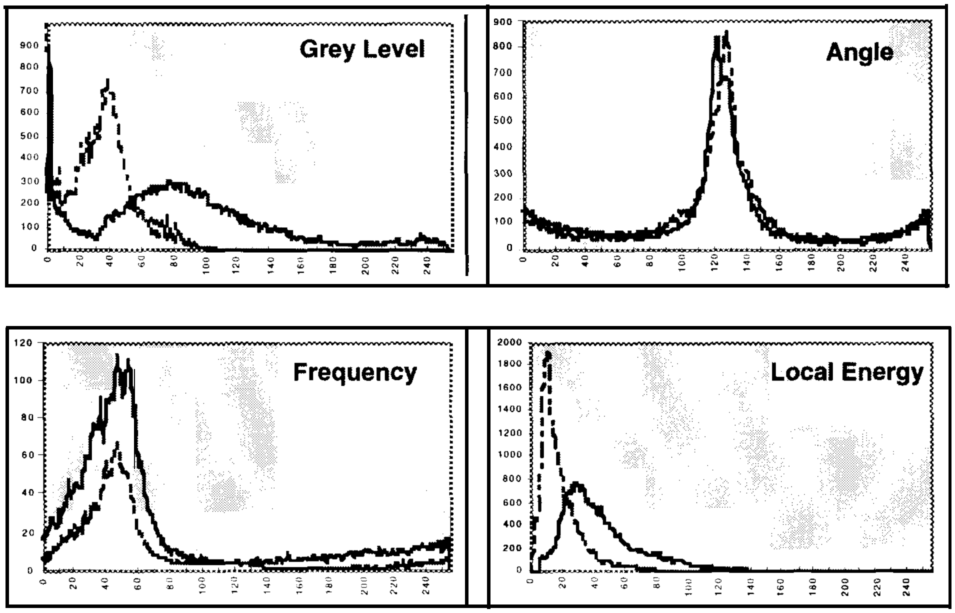

— Real — **Synthetic**

Figure 5. CAMAELEON histograms.

A wide dynamic range of signature image pairs is created from field measured data and its corresponding synthetic scene or from field data with differences in either time, background, target, environment, or other measurable variables. Each image set is parsed through a comparator using, in this case, the German computational vision model CAMAELEON [4, 5, 9]. This model provides, as an output, the first-order metrics such as grey level, frequency, orientation, and local energy distributions. For each of these metrics, a statistical analysis is applied in terms of mean, median, mode, variance, standard deviation, absolute deviation, skew, kurtosis, and entropy. Figure 4a shows the field-measured data of an M60-A1 tank scene taken with a calibrated infrared sensor at Fort AP Hill, VA. A data artifact was introduced unintentionally because every other field was missed during the digitization process, resulting in a lower quality image than what a high-resolution, calibrated DL FLIR [16] is capable of showing. Figure 4b is the synthetic infrared rendering by the CREATION model. It contains some statistical sampling rather than first principle rendering of background data. Figure 5 shows the comparison between the real and the synthetic scene based on the output of the CAMAELEON model. Figure 6 shows our use of the similarity metric with our present validation approach. Two noncomputational vision approaches developed by ARL for comparing images, the "Region-Based" and the "Symmetric Difference" methods, are discussed next.

Figure 6. Image similarity.

3.1 A Region-Based Method of Comparing Images

Historically, there have been two different approaches for comparing a real scene to the synthetically rendered scene. One relies on very local information such as pixels, and the other relies on global information such as statistics generated over the entire image. Our approach is to use the middle ground between these approaches. We call it a *region-based* method of comparing images.

In region-based image comparison, we produce a mask that defines nominally 40 to 100 regions in the real or the synthetic image under comparison. This mask is then applied to both the real image and the synthetic image, and image-statistics (grey-level histograms, local-energy, etc.) are computed over each region. Comparisons are then made from the statistics obtained from each region to the statistics obtained from the same region in the other image. A real-synthetic image comparison is then obtained by computing the area-weighted average of the similarity metrics that resulted from each of these local comparisons.

For example, note that the images in Figure 2 have identical global grey-level histograms. Thus any comparisons based on global grey-level histograms will indicate that the two images compare identically. Suppose, . alternatively, that we generated any arbitrary tessellation and applied that tessellation to both images. We could then generate the grey-level histograms over each of the regions produced, and then compare the histogram of a particular region to the histogram in the corresponding region in the other image. While the grey-level statistics over the images as a whole were identical, the same statistics over a small portion of each of the images would most likely not be.

Thus, regionalization has the potential to reduce the negative effects of false matches that can occur when comparisons based on global image statistics are used. Additionally, the use of regionalization cannot produce results that make images appear to be statistically more similar than that which would be obtained from a global comparison. Specifically, if the images compare favorably on a region-to-region basis, they will also compare favorably globally. In short, regionalization can be used to allow a more rigorous application of common global statistics, and achieves a balanced medium between the very demanding bit-wise image comparisons and the somewhat ineffectual global comparisons.

As to the question of *which* regionalization to use, it is the opinion of the authors that any regionalization or tessellation would be acceptable since an arbitrary regionalization has the potential for producing more accurate image comparisons than the corresponding global comparison, and no regionalization has the potential for producing comparisons that are less appropriate than global comparisons. We, however, also see an advantage in producing regionalizations that are consistent with the low-frequency variations that are naturally present in the image.

3.1.1 A Method of Comparing Images Under a Low-Frequency Mask

In this section we describe step-by-step a region-based approach of comparing two images. Step 1 aligns the two images with each other, Step 2 matches the average brightness of one image to that of the other image, Step 3 creates a low-frequency mask, and Step 4 uses the mask to apply a global image comparison metric.

Step 1: Register and crop the images.

The first step in comparing images is to make sure they are (1) registered (or aligned) with each other and (2) of identical size. We accomplish this as follows. A human operator locates a well-defined point in each image and determines the pixel coordinates of that point. At present, a tool such as *xv* (x-windows view) is used to determine the coordinates of that point. Note that the real and synthetic images have not been previously registered, so the coordinates of the "well-defined point" will in general be different for the two images. Our program *crop* is then used to produce a new image for each of the two original images. The new images will be of a specific size and centered at the chosen point. We used a size of 256 pixels by 256 pixels for the initial trials. At this point we have a real-synthetic pair of images that are the same size and registered with each other.

This registration process involves translation only. A more sophisticated method of registration would also include rotation and scaling. However, this would necessitate a method of mapping the rectilinear grid of pixels of the original image to another grid of a different angular orientation and having a different scale. Methods for performing such a mapping exist, though they generally introduce artifacts, and the artifacts would have a great potential to pose problems for the comparison. Thus, we insist that the real and synthetic images be made to the same scale and oriented at the same angle, and then the registration is done by performing only a translation.

Step 2: Gamma-match the images.

A process that we call gamma-matching is used to change the brightness levels of the synthetic image so that the total brightness of the synthetic image is within a small factor of the total brightness of the real image. This is done with gamma-correction as follows. If the synthetic image is dimmer than the real image, gamma values are successively chosen starting at 1 and increasing in steps of 1 until the brightness of the corrected synthetic image exceeds the brightness of the real image. If the synthetic image is brighter than the real image, gamma values are successively chosen starting at 1 and decreasing in steps of 0.1 until the brightness of the corrected synthetic image is less than the brightness of the real image. When such cross-over points are found, the process is repeated with smaller step sizes, specifically step sizes of one tenth the size currently being used. This entire process is repeated until the ratio of the total brightness of the synthetic image to the total brightness of the real image is within a small constant of 1. We have been using the constant 0.001.

Step 3: Construct a template for one of the two images.

The process of constructing a low-frequency template for an image involves three steps: low-pass filtering the image, multilevel thresholding the image, and uniquely labeling the regions that result. These three steps are explained next.

Step 3a: Low-pass filter the image.

The image is low-pass filtered by convolving it with a square template containing a normalized Gaussian function.

Here, the user specifies the size of the "radius" of the template to be used and then the template is automatically generated. For example, if the template radius is chosen to be 5, then an 11 by 11 template will be generated. This template is then filled with a two-dimensional Gaussian function centered at the center pixel in the template and having a sigma chosen so that the area under the Gaussian curve over the template is 99% of the total area under the Gaussian curve over the entire real plane. The entries in the template are then normalized over the template (i.e., scaled so that they sum to 1). When the image has been convolved with such a template, the brightness values of the image will be "smooth" in the sense that there will be no rapid changes in brightness. The image will actually appear blurred. This is done so that the next step of thresholding will not produce many small regions, which would be possible if the image contained high-frequency components.

Step 3b: Apply multilevel thresholds to the image.
After Step 3a, the image will have brightness values that do not change rapidly throughout the image. In fact, the image can be thought of as a landscape in which the brightness values represent altitudes. After the low-pass filtering, all the hills and valleys will be smooth and gently rounded. There will be no sharp peaks, no spikes, no cliffs, no abrupt changes of any sort.

We now apply a multilevel thresholding process to the low-pass filtered image. Here, we (arbitrarily) choose 4 levels of brightness uniformly spaced between the dimmest pixel value and the brightest pixel value. Thus, every single pixel in the image falls into exactly one range. Each pixel is then assigned a number from 1 to 5 corresponding to the range into which it falls. After this thresholding process, the image appears somewhat like a topographic map in that the regions shown correspond to the brightness range of the pixels in the region. This looks very similar to topographic maps that show the altitude ranges between the curves on the map.

Step 3c: Label the regions produced.
The next step is to give the region that result from the above step a unique label and to give all pixels the label of the region into which they fall. To accomplish this, we initially set all pixels to be unlabeled and then start in the upper left comer of the image and proceed in a raster-like scan. During the scan, we do the following. When an unlabeled pixel is encountered, we use a recursive process that we call "flooding" to label each pixel in the same region as the newly encountered pixel with the "next available label". When the single raster scan completes. all pixels in the image will be labeled, either from the scan itself or from the recursive flooding process.

Here we scan the image from left to right and from top to bottom in a raster pattern. If a pixel is encountered that is unlabeled, it is given the "next available label" and the recursive "flooding" procedure is invoked. This procedure will attempt to recurse one pixel left, one pixel right, one pixel up, and one pixel down from the recently labeled pixel. Specifically it will recurse in those directions if and only if the pixel in that direction has not been previously labeled, the pixel in that direction exists (i.e., it is not at the edge of the image), and the pixel in that direction has the same brightness value (after the thresholding) as the recently labeled pixel. Thus the recursion will proceed to label all pixels in the given region and it will not proceed across region boundaries. At the completion of recursively labeling a region, the raster scan will continue until another unlabeled pixel is encountered at which point the region containing it will be flooded in a similar manner. This continues until the entire image has been scanned and labeled.

Step 4: Apply a global-metric according to the mask produced.
Upon completion of the mask, or tessellation, that was produced by Step 3, we can now use that mask to apply a similarity metric in a region-based fashion. As an example, let us consider using the "common area overlap" of the gray-level histogram as our metric. Recall that if we were using this metric in the global sense, we would generate the gray-level histogram for the real image and also for the synthetic image, normalize each of those histograms, and then determine the area under the histograms that is common to both histograms. Note that this area will be zero if the histograms are completely disjoint and it will be one if the histograms are identical. In the region-based method, we will apply this metric not to the images as a whole, but to each of the individual regions of the tessellation. Thus we will obtain a number between zero and one for each of the regions that indicates the extent to which the gray-level histograms of the two images resemble each other in that region. To obtain the final image metric, we simply take an area-weighted average of the similarity metrics that we obtained for the regions.

3.1.2 The Symmetric Difference Method

The segmentation of images into regions as described above can also be used to produce a figure of metric of comparison between two images. Consider generating one mask based on the real image and another mask based on the synthetic image. Comparing the masks, then, could be used as a method of comparing the images.

The manner in which we compare the masks is as follows. For one of the two masks, say that derived from the real image, consider a particular region in the mask. Calculate the symmetric difference between this region and every region in the mask of the other (synthetic) image and choose the minimum. Notice that if this region closely coincides with a region in the other image, this minimum will be small. In fact, if there is a region in the other mask that is identical to the region under comparison, this minimum will be zero. Now choose a second region in the real image and repeat this process to obtain another “minimum symmetric difference”. If we continue this for all regions in the image, we will develop a string of minimum symmetric differences. We use the sum of this sequence as the image metric.

Notice that this sum will be small when each region in the real image closely coincides with a related region in the synthetic image, and it will be large otherwise. Thus, we have a metric that will be zero upon comparing two identical images, it will be small when comparing two images with similar low-frequency structures, and it will increase for images that have few such commonalties.

4. GOALS

The near-term goals are to be able to generate a sufficient number of image pairs over a wide dynamic signature range based on high-quality field measurement data. Investigation of other candidate validation metrics [19] will need to be analyzed.

The long-term goal is to compare various target acquisition models in terms of their capability to predict second-order effects, again over a wide dynamic signature range. The prediction results from these various models can be compared to the baseline (human observers), and the differences can be analyzed using validation metrics identified to date from this work.

A vision to allow optimization of resources in this particular research area is to apply this methodology for dual technology application (military and nonmilitary). The methodology being developed for validation will lend itself well for the analysis to allow improvement of target acquisition modeling. One pristine area for target acquisition enhancements for the military, as well as for law enforcement, is in the littoral environment [12, 13, 15]. A coalition force is being attempted between DoD agencies that could be extended to NATO working groups that are interested in this particular area of R&D.

5. ACKNOWLEDGMENT

This effort was made possible by the members of the ARL Synthetic Imagery Branch [Khang Bui, Janice Colby, Glenn Dockery, Giap Huynh, Teresa Kipp (Branch Chief), Hung Nguyen, Joseph Penn, and Paul Zirkle], who provided the real-synthetic image pairs presented in this paper, and who provided many hours of consultation and technical support. Further technical support was provided by Matthew Thielke of the ARL Automatic Target Recognition Branch. Jim McManamey and Eddie Jacobs, members of the Night Vision Electro-Optics Laboratory Camouflage Modeling Branch, Fort Belvoir, Virginia, ran our images through the CAMAELEON software package. This includes the images presented in this paper and many, many others.

[20] *Gary* Witus, and Thomas Meitzler, *TARDEC Visual Perception Model, Calibration and Validation,* Seventh Annual Ground Target Modeling and Validation Conference, August 20-22, 1996, Warren, Michigan

[21] Andrew Zembower and Marcos C. Sola, *Signature Quality Metrics,* Eleventh Annual Ground Vehicle Symposium, KRC, 22 August 1989.

7.0 Biographical Sketches of the Authors

Mr. Marcos C. Sola is the technical leader of the Modeling Validation Group in the Synthetic Image Branch of the U.S. Army Research Laboratory in Adelphi, Maryland. His research interests include synthetic signal validation, mathematical modeling of target signatures, modeling of atmospheric effects on multispectral electromagnetic wave propagation, and systems performance analysis. In addition to his current position, Mr. Sola's career has included work as Technology Associate for Signatures for the Sensors, Signatures, and Signal Processing Office, Project Leader for the Survivability and Management Office, the CM/CCM Office, and the Night Vision and Electro-optics Laboratory in Fort Belvoir Virginia. Mr. Sola holds B.S. and M.A. degrees in Physics from American University, has received numerous U.S. Army Achievement Awards. and publishes extensively.

Dr. Mark Walter Orletsky was born in Johnstown, Pennsylvania on September 25, 1960. He reeeived a Bachelor of Science degree in Electrical Engineering from the Pennsylvania State University in 1982, a Master of Science degree in Electrical Engineering and Computer Science from the Johns Hopkins University in 1984, and the degree of Doctor of Philosophy in Computer Science from the Johns Hopkins University in 1996. Dr. Orletsky's research interests include Computational Geometry, Computer Graphics, Fault-tolerant Computing, and Computer Vision. He currently holds a position with the U.S. Army Research Laboratory in Adelphi, Maryland where his efforts are focused toward multispectral synthetic scene production and validation.

Mr. Quochien B. Vuong received B.S. and M.S. degrees in electrical engineering from George Masson University in 1988 and 1993, respectively. He was a research engineer at the U.S. Army, CECOM, Center for Night Vision and Electro-optics, Fort Belvoir, Virginia from 1990 through 1992. During this period Mr. Vuong worked primarily in the area of optical signal processing and had received two U.S. Army Achievement Awards. In 1992, Mr. Vuong joined the U.S. Army Research Laboratory in Fort Belvoir, Virginia. He is currently a research engineer at the U.S. Army Research Laboratory in Adelphi, Maryland, where his present research interests include digital image processing, and the production and validation of synthetic imagery.

Mr. Charles R. Kohler was born in New York City and received a Master of Science degree from Cornell University, Ithaca, New York in Engineering Physics and Mathematics. He was commissioned a 2ND Lieutenant, ROTC, Infantry, in the US Army and spent about 3 years on active duty in Germany. He is "Airborne" and "Special Forces" Qualified, a Vietnam Era Veteran, and currently a member of the US Army Ready Reserves. He has over 38 years of federal service and was recently selected by the Department of the Army to be a Judge at the International Science and Engineering Fair. He has a keen interest in simulating science and engineering, education, among American youth. He has been an outstanding volunteer judge at many science fairs in the past. His current interest is modeling, simulation, and synthetic imagery.

6. REFERENCES

[1] *E. L. Ayers, J. F. Nicoll, Comparison Between IRTOOL-Simulated Ocean-Sky Scenes and IRAMMP Data for Target Detection,* IDA Task #A- 180, ARPA/STO.

[2] T. J. Doll, S. W. McWhorter, D. E. Schmieder, *Computational Model of Human Visual Search and detection.* Proceedings of the March 1994, IRIS Passive Sensors Symposium, Albuquerque, New Mexico.

[3] T. J. Doll, S. W. McWhorter, D. E. Schmieder, A. A. Wasilewski, and G. Welch, *Georgia Tech Vision (GTV) Model,* Version GTV94a, Analyst's Manual, Georgia Tech Research Institute., Atlanta, Georgia.

[4] R. Hecker, *CAMAELEON - Camouflage Assessment by Evaluation of Local Energy, Spatial Frequency, and Orientation,* Proceedings of the Third Annual Ground Target Modeling and Validation Conference, Anne Marie L. LaHaie, Ed, U.S. Army TACOM, Warren MI (1993).

[5] R. Hecker, "CAMAELEON - Camouflage Assessment by Evaluation of Local Energy, Spatial Frequency, and Orientation", *Characterization, Propagation, and Simulation of Sources and Backgrounds 11,* Dieter Clement and Wendell R. Watkins, Editors, Proceedings of SPIE- 1687, Pages 342-349.

[6] Charles Kohler, and Marcos Sola, *CREATION - Multispectral Synthetic Scene Generation,* Seventh Annual Ground Target Modeling and Validation Conference, August 20-22, 1996, Warren, Michigan.

[7] G. H. Kornfeld, *Various FLIR Sensor Effects Applied to Synthetic Thermal Imagery*, Proceedings of SPIE, April 1993.

[8] G. H. Lindquist, G. Witus, T. H. Cook, J. R. Freeling, and Grant Gerhart, *Target Discrimination Using Computational Vision Human Perception Models,* Proceedings of the SPIE Meeting, Arpil 7, 1994.

[9] James R. McManamey, *An Initial Investigation of the CAMAELEON Model,* Proceedings of the Sixth Annual Ground Target Modeling and Validation Conference, August 1995, Houghton, Michigan, Pages 239-248.

[10] Joseph A. Penn, Hung Nguyen, Teresa Kipp, Charles Kohler, Giap Hyunh, and Marcos Sola, *The CREATION Scene Modeling Package Applied to Multispectral Missile Seekers and Sensors,* Proceedings of the 1995 Conference on Multispectral Missile Seekers and Sensors, November 1-2, 1995, Redstone Arsenal, Alabama.

[11] Joseph A. Penn, H. Nguyen, M. Sol a, C. Kohler, J. Weber, and S. Hawley, *The CREATION Scene Modeling Package Applied to Theater Air Defense Fire Control Situations,* Proceedings of the 1995 National Fire Control Symposium, Montery, California, July 31- August 3, 1995.

[12] Marcos C. Sola, *Active/Passive Signature Enhancer,* Patent Application, ARL Docket # 96-34, 26 October 1996.

[13] Marcos C. Sola, *Application of Optical Augmentation Technique for Maritime Search and Rescue Operations,* National Target/Threat Signature Data Systems Conference, Naval Air Warfare Center, Point Mugu, California, 23-25 July 1996.

[14] Marcos C. Sola, and Tadeusz M. Drzewiecki, *A Proposed Technical Methodology to Reduce the Number of Physical Infrared Signature Measurements,* Eleventh Annual Ground Vehicle Symposium, KRC, 22 August 1989.

[15] Marcos C. Sola, Joseph A. Penn, E. Glenn Dockery, Paul Zirkle, Charles R. Kohler, Teresa Kipp, and Janice F. Colby, *A Proposal for an Active/Passive Signature Enhancer (APSE) for IFF*, Accepted for publication in 1997 Joint Service Combat Identification Systems Conference (CISC-97), 15-17 April 1997, Navy Amphibious Base Coronado, San Diego, California.

[16] William Stump, Notes from *M60A1 Tank Diurnal IR Signatures, 1995* Night Vision and Electro-optics Division (NVESD) Report.

[17] J. Weber and J. A. Penn, *CREATION User's Manual* (draft) Version 1.27, U.S. Army Research Laboratory, Signature Modeling Branch, Fort Belvoir, Virginia, S3ID/SMB, November 1994.

[18] J. Weber and J. Penn, *Creation and Rendering of Realistic Trees,* SIGGRAPH 95 Computer Graphics Conference Proceedings, Annual Conference Series 1995, pp. 119-128, ACM SIGGRAPH, August 1995.

[19] James P. Welsh, *Smart Weapons Operability Enhancement (SWOE), Joint Test and Evaluation (JT&E) Program Final Report,* SWOE Report 94-10, August 1994.